Sharing the Power of Supercomputers

More researchers are getting ready access to Livermore's powerful supercomputers for advanced simulations.

COMPUTER simulation is changing the nature of scientific discovery by becoming a full partner with theory and experiment. Nowhere is that transformation more visible than at Lawrence Livermore, where scientists are relying increasingly on computer simulations, especially those based on high-resolution, three-dimensional models. Such simulations, however, require the enormous computational horsepower of the latest generation of supercomputers.

To ensure that all Livermore programs and researchers have the possibility of accessing high-performance computers, the Laboratory created the Multiprogrammatic and Institutional Computing (M&IC) Initiative in 1996. The initiative has grown substantially in the last few years and currently serves more than a thousand users, including outside collaborators, with some of the most powerful computers available. Indeed, virtually every Livermore unclassified program—from atmospheric sciences to bioscience—benefits from the computing resources provided by the initiative.

"M&IC is a partnership between the Laboratory and its research programs to bring high-performance computing to every researcher," says Mike McCoy, head of Livermore's Scientific Computing and Communications Department. M&IC recognizes that no matter what their mission, scientists should have ample access to centralized, high-performance computers far more capable than the computational resources that individual departments could afford to purchase and support. M&IC leaders, under the direction of McCoy and former Computation Associate Director Dave Cooper, have worked with key researchers around the Laboratory to create a centralized operation that helps all researchers do better science through more advanced simulations.

McCoy notes that his department's goal is to make the "discovery environment" that is evolving in the Accelerated Strategic Computing Initiative (ASCI) available to all researchers. ASCI is an element of the National Nuclear Security Administration's Stockpile Stewardship Program to assure the safety and reliability of the nation's nuclear weapons in the absence of underground testing. (See *S&TR*, June 2000, pp. 4–14.) In the ASCI computing approach, scientists are supported by powerful simulation tools that make possible both the generation of raw scientific data and the manipulation of this information to enhance understanding.

The institutional nature of the initiative makes it possible for anyone with a good idea to perform teraops-scale (more than 1 trillion mathematical operations per second) parallel computing to achieve breakthrough scientific discoveries. First developed in the late 1980s, parallel computing attacks huge mathematical problems with a number of identical (and typically inexpensive and common) processors simultaneously sharing a computational task. Parallel computing differs from traditional supercomputing, in which a few high-cost, specially designed vector processors perform numerical calculations.

"Parallel computing brings formerly insoluble problems into solution range and makes formerly difficult-to-solve problems routine," McCoy says. Examples of parallel supercomputers include the Laboratory's ASCI supercomputers manufactured by IBM—the classified Blue Pacific (3.9 teraops) and White (12.3 teraops) and the unclassified Blue (0.7 teraops) and Frost (1.6 teraops)—and M&IC's

TeraCluster2000 supercomputer, built by Compaq, operating at 0.7 teraops.

McCoy explains that when the initiative was launched in the mid-1990s, many scientists, having experienced difficulties in securing sufficient and sustained access to unclassified centralized computing, had forsaken high-performance computing entirely. Instead, they had taken advantage of the opportunities presented by new and relatively inexpensive and powerful desktop computers in their offices or used terminals tied to scientific workstations owned by their research programs. He calls the situation at the time "a desktop diaspora."

While desktop workstations continue to be an important tool for scientific computing research, these machines do not always provide the necessary computational power. "Exclusive recourse to workstations, in the absence of access to the most powerful centralized systems available, could have left many Lawrence Livermore researchers unequipped with the computer tools they needed to remain competitive in the next decade," says McCoy.

"Until this time, most scientists were performing two-dimensional simulations because three-dimensional simulation was quite difficult, if not impossible, with the computers at their disposal," says Greg Tomaschke, M&IC computing systems leader. "Livermore has a long history of preeminence in computational simulation, and people were concerned they'd start to lose this advantage." M&IC has ensured that

On the Horizon: The Terascale Simulation Facility

Looking to the future, Lawrence Livermore managers are planning a new facility to house their newest supercomputers. Called the Terascale Simulation Facility, the building will cover 25,000 square meters and have offices for 288 people. A two-story main computer structure will provide the space, power, and cooling capabilities to simultaneously support two future supercomputer systems for the National Nuclear Security Administration's Accelerated Strategic Computing Initiative (ASCI). Nearby, the Livermore Computing Center will house the unclassified computing resources of the Multiprogrammatic and Institutional Computing Initiative.

The computer structure will be flanked on the south side by a four-story office building. The building will have space for Lawrence Livermore staff, vendors, and collaborators from the ASCI University Alliance and visiting scientists from Los Alamos and Sandia national laboratories.

Lawrence Livermore computer managers say that no existing computer facility at Livermore is adequate, nor can any be modified sufficiently, to accommodate the newest supercomputers, which are the largest computational platforms ever built. These platforms are so powerful—exceeding 50 teraops, or 50 trillion operations per second—that they are increasingly labeled ultracomputers. Current plans are to site a 60-teraops system in the Terascale Simulation Facility by 2004.

The Terascale Simulation Facility will house the largest supercomputers ever built and offices for nearly 300 people.

Livermore researchers are at the forefront of simulation science, he says.

Serving Programs and Individuals

The M&IC is so named because it serves both research programs (multiprogrammatic) and individual (institutional) researchers. A research program can either purchase a block of time on existing machines or, more popularly, share in the investment in new equipment, which gives the program access proportional to its investment.

The research program designates a principal investigator to select users of the resources it has paid for. The Livermore Computing Center furnishes the principal investigators with Internet-based records of the computing time the program is entitled to and how much of the time has been used, in total or by individual researcher. A hotline is available to answer technical computing questions and provide account information.

M&IC also grants computer time to individual researchers, independent of their connection to a program, whose work is often viewed as less mainstream than most efforts in their particular research specialties. The researchers are selected through a short proposal process by the Institutional Computing Executive Group (ICEG), composed of users from the Laboratory's research directorates. The ICEG provides general oversight of M&IC, in cooperation with the Computation associate director and the Laboratory's deputy director for Science and Technology.

"The ICEG is the most important Livermore Computer Center link into the user community," says McCoy. "The M&IC Initiative is based on the strong bonds of support and advice that exist between the ICEG and the center. As a result, Livermore supercomputing has become an institutional resource much like the library, a place where researchers from any program can expect resources to support their research."

Researchers can monitor the status of their simulation from their desktop computer. Those who sign up for long simulation runs can see their place in the queue. Once the simulation begins, they can open a window on their computer screen to watch the calculations. "We've made investments in customer (user) support that significantly exceed the norm for computing centers worldwide because we know the extraordinary challenges scientists face with writing and running large parallel calculations," says McCoy.

Providing Capacity and Capability

Under the M&IC Initiative, the Livermore Computing Center has acquired increasingly more powerful clusters or groups of computers. The computers are of two conceptually different kinds: capacity computers and capability computers. Capacity computing is designed to handle jobs that don't require a lot of computing horsepower or memory. It is often used for quick turnaround of a large number of small to moderately sized simulations.

Capability computing uses a substantial fraction of the entire computing power of a supercomputer to address a large-scale scientific simulation in three dimensions. "The driver for capability computing usually is the need for large amounts of memory, which means harnessing many processors to work together," says Tomaschke. A capability computing resource can only serve a few users simultaneously.

McCoy says that any effective computational environment is supported by a capacity foundation. "Capacity allows users to develop the applications and work the studies that are necessary to conceive of, develop, and debug capability applications," he says. M&IC managers, working with the ICEG, developed a strategy early on to first build a capacity foundation and keep it current. They then devised a second strategy to build access to capability computing, either through a partnership with the ASCI program or through unique research and development relationships with major vendors. Both strategies have proved effective.

The Compaq 8400 Compass Cluster was the first capacity computer resource sponsored by the M&IC Initiative. Delivered in 1996, the Compass Cluster consists of 8 nodes (rack-mounted computers), with each node possessing between 8 and 12 central processing units (CPUs, or microprocessors). In all, 80 CPUs provide 7 gigaops (7 billion calculations per second) of computing power, 56 gigabytes of memory, and about 900 gigabytes of disk space. A replacement for Compass, scheduled to arrive late this year, will provide a total of 192 gigaops.

To increase capacity computing, M&IC acquired TeraCluster in late 1998. In all, TeraCluster consists of 160 CPUs, 80 gigabytes of memory, and over 1.5 terabytes (trillion bytes) of disk space and provides about 182 gigaops of computing power. It is closely integrated with the Compass Cluster.

In September 2000, the Livermore Computing Center took delivery of the Linux Cluster, which generates 42 gigaops of computing power. It is composed of 16 advanced Compaq nodes, with each node having 2 CPUs and 2 gigabytes of memory. The machine is used to increase computing capacity and provide users with an opportunity to evaluate the potential advantages of the Linux operating system. McCoy says that Linux represents a radical departure because it is not a proprietary operating system. Success with the machine could lead to the procurement, next year, of Linux parallel processing systems using a high-performance interconnect and advanced Intel processors.

The M&IC Initiative has also acquired a system manufactured by Sun Microsystems called Sunbert, which provides 12 gigaops of power with 24 CPUs, 16 gigabytes of memory, and approximately 600 gigabytes of disk space. The system is designed to allow access by Livermore researchers who are foreign nationals from sensitive nations.

Capability Power

The M&IC's most important capability platform, the 680-gigaops TeraCluster2000 (TC2K) parallel supercomputer, arrived last year. The machine was the result of a three-year Cooperative Research and Development Agreement between computer scientists at Livermore and Compaq Computer Corporation to evaluate a new supercomputer design based on Compaq's 64-bit Alpha microprocessor and Quadrics Corporation's interconnects and software. The alliance resulted in the Compaq SC Series of supercomputers, of which TC2K is serial number 1.

Tomaschke says that the most important aspect of Livermore's role in the collaboration with Compaq was providing advice based on many years of experience with supercomputers. "We gave Compaq important feedback about what scientists require for doing three-dimensional simulations, such as operating enormous file systems."

The TC2K consists of 128 nodes, with 4 Alpha processors per node. In total, the machine has 512 CPUs, 256 gigabytes of memory, and 10 terabytes of disk space. The 128 nodes are partitioned like a giant hard disk. The largest partition is dedicated to the most complex simulations, while a small partition permits a researcher to interact with the machine in real time. Occasionally, all nodes are freed up for a single task, such as experiments to determine if a code will scale properly when the number of nodes increases sharply.

Limited availability of TC2K began early this year, with 25 projects (involving about 100 researchers) "shaking down" the machine. It became generally available in August, vastly increasing computing capability to all unclassified researchers.

TC2K represents one of three capability resources. The second resource is ASCI Blue, the 740-gigaops unclassified portion of the ASCI Blue Pacific system, which has 282 nodes with 4 IBM Power PC processors each. The third resource is ASCI Frost, the unclassified version of ASCI White. This system features 68 nodes with 16 powerful 1.5-gigaops IBM processors each and 16 gigabytes of total memory. This computer peaks at 1.6 teraops and is both the most modern and most powerful unclassified computer on site. Although both Blue and Frost are primarily dedicated to the ASCI mission, significant access has been made available to a number of Livermore science teams.

TC2K's capability, combined with that of ASCI Blue and ASCI Frost, provides unprecedented unclassified computing capability for a national laboratory, says McCoy. Researchers perform code development and limited simulations on capacity machines, complex three-dimensional simulations

The TeraCluster2000 (TC2K) Compaq supercomputer, operating at 0.7 teraops (trillion operations per second), has 512 central processing units arranged in 128 nodes, 256 gigabytes of memory, and 10 terabytes of disk space.

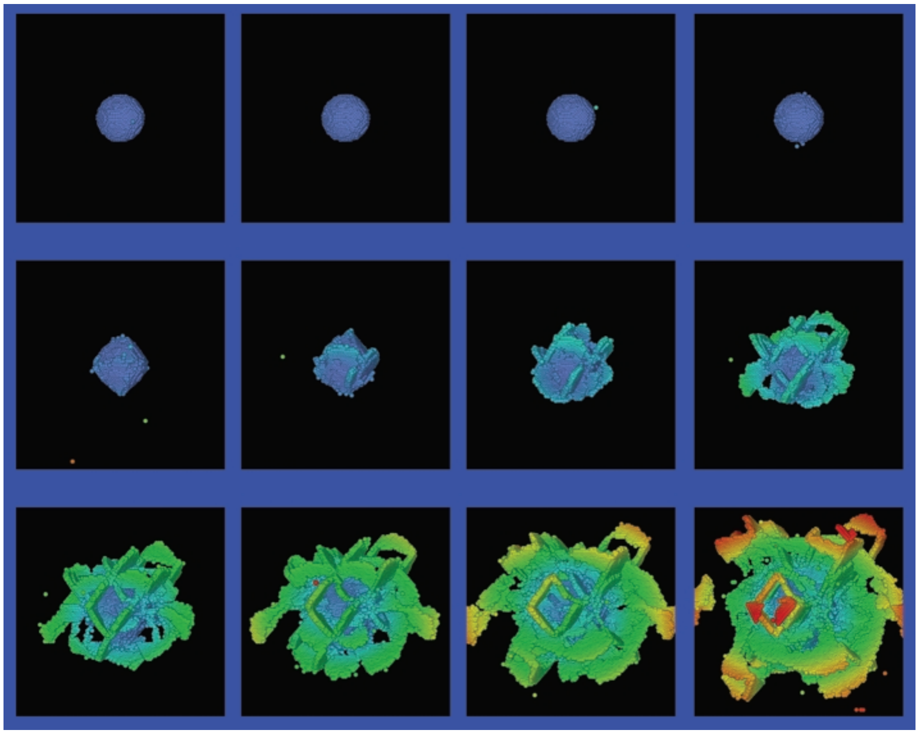

A sequence of snapshots from a simulation showing how a growing fracture deforms a stressed copper crystal at the level of atoms. The first row shows no defects. The second row shows defects forming on the surface. The third row shows defects shooting off from the surface in a process that forms a large hole.

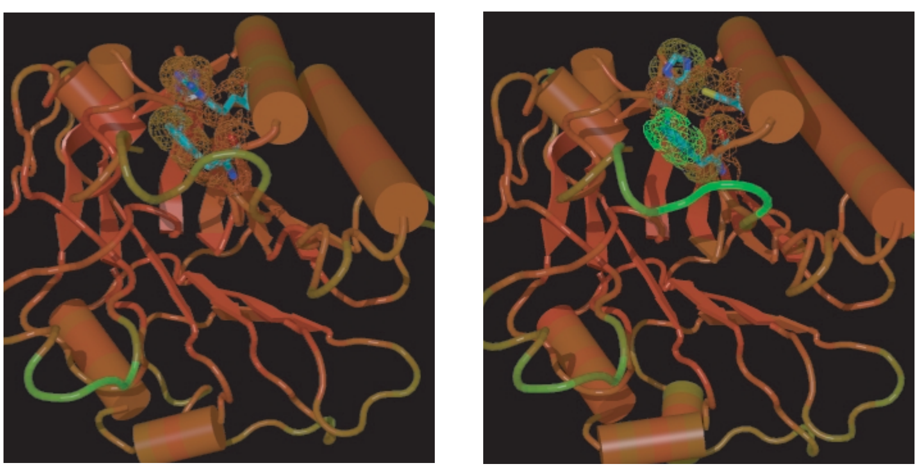

The TC2K supercomputer is used to gain insight into the human Ape1 enzyme, a protein that repairs DNA. The simulation at left is a healthy protein. The simulation at right is a version of the protein that contains a single amino acid substitution. This variant shows much more motion in the front loop of the protein; the motion is a means of recognizing DNA damage. The coloring indicates the amount of intramolecular motion, with reddish-brown being the least motion and greenish-blue the most motion.

on TC2K and Blue, and the most demanding runs on Frost. Access to the limited unclassified ASCI resources is extremely competitive.

Never Before Attempted

The resources provided by the M&IC Initiative are permitting researchers to generate simulations that, in many cases, were never before attempted for lack of computing power. As a result, the Laboratory is at the forefront of simulating a wide range of physical phenomena, including the fundamental properties of materials, complex environmental processes, biological systems, and the evolution of stars and galaxies.

For example, physicist Burkhard Militzer of the Quantum Simulation Group in the Physics and Advanced Technologies Directorate is using TC2K to simulate how gases such as hydrogen and oxygen behave under extreme pressure and to compare those simulations with results of gas-gun experiments done at Livermore. Militzer uses JEEP, a parallel supercomputer code developed by Livermore physicist Francois Gygi. The simulations typically require upward of 300 hours of processing time. Because there is a time limit of 12 continuous hours on TC2K, the simulations run in chunks.

Militzer says that he would like to use JEEP's quantum mechanics capability to simulate the weak hydrogen bonds that keep two DNA strands in their helix. "It's very difficult to do accurately because of all the water molecules surrounding the DNA," he says. "TC2K enables a new class of projects. I wouldn't even begin to think about running a simulation of DNA hydrogen bonding without having TC2K."

Examining the Birth of Cracks

Physicists Robert Rudd and Jim Belak run simulations on the Compass

Cluster. The simulations examine in microscopic detail the birth of fractures in metals such as copper under the extreme stresses of a shock wave. The molecular dynamics simulations are done in a nanometer-scale box holding about one million virtual atoms.

The simulations, actually a sequence of snapshots from a movie, depict a copper crystal that is deformed by the growing fracture over a period of 60 picoseconds (trillionths of a second). Only the atoms at the fracture surface or in crystal defects are shown. The defects, known as dislocations, can be seen shooting off in a process that forms an increasingly large hole or fracture. Rudd says, "We're interested in learning more about how voids grow and how the material deforms around them."

"We have great confidence in our simulations," adds Rudd. He is planning to use ASCI Blue to vastly expand the length of the simulated piece of metal and to simulate much longer time periods.

Lawrence Livermore biological scientists in the Computational Biology Group have been one of the most visible users of the M&IC Initiative. (See *S&TR*, April 2001, pp. 4–11.) The researchers have produced stunning depictions of DNA and proteins that reveal the exact mechanisms of key biological processes. Parallel supercomputers are ideal for this kind of simulation because they excel at modeling the interactions of large numbers of atoms contained within biological macromolecules.

Researcher Daniel Barsky of the Computational Biology Group has been studying the dynamics of Ape1, an enzyme responsible for repairing a common form of DNA damage called abasic lesions. Barsky uses the TC2K to compare the degree of intramolecular motion of normal Ape1 with a variant found in which the enzyme contains a single amino acid substitution.

Computer scientist and geophysicist Shawn Larsen has been using TC2K to do three-dimensional simulations of oil exploration problems and seismic wave propagation. Larsen's E3D code is optimized for parallel supercomputers. (See *S&TR*, November 2000, pp. 7–8.) His oil exploration modeling entails so-called elastic simulations that are about 250 times more computationally intensive than standard acoustic (sound waves in air) simulations. Elastic simulations provide better details of subsurface geology, which is essential to oil exploration efforts, by depicting both the S (shear) waves and the P (compressional) waves that travel in water and earth. With computers such as TC2K, realistic three-dimensional elastic simulations are now possible.

Climate Simulations Set the Pace

Livermore researchers in atmospheric sciences were among the first to take advantage of M&IC's new capabilities. Having performed parallel computing for 10 years, they had already achieved some of the most advanced simulations ever done. They had also made their simulation codes portable, that is, easily adaptable to different computers. As a result, it did not take long to adapt their codes to TC2K and ASCI Frost.

They put TC2K to the test to study the effectiveness of ocean carbon sequestration, a proposed approach for mitigating global warming. In one sequestration model, the carbon dioxide generated by industrial operations would be injected into the oceans instead of being emitted into the

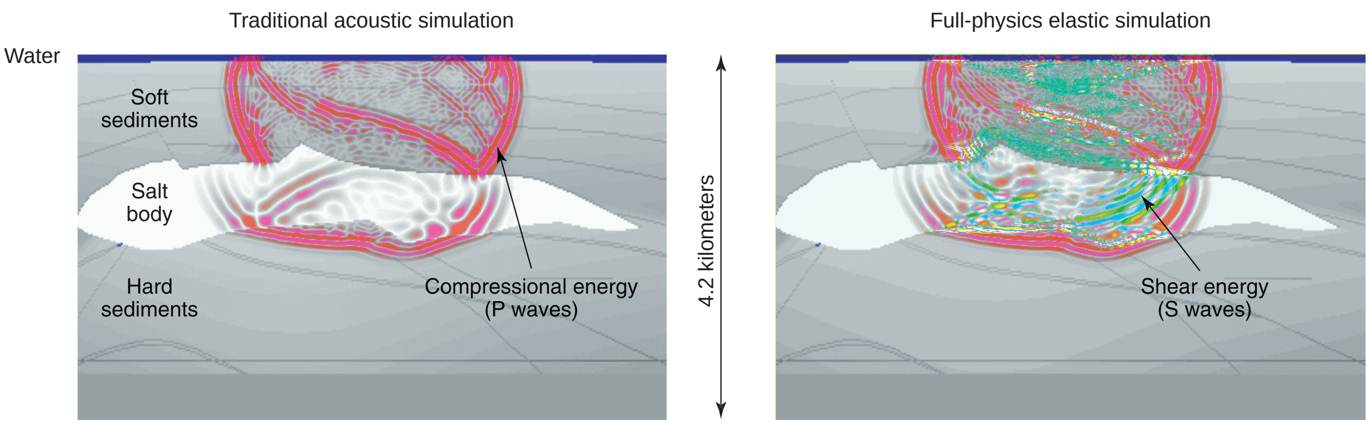

The TC2K supercomputer is able to perform three-dimensional simulations of seismic wave propagation with Livermore's E3D code that is optimized for parallel supercomputers. The image on the left used 1 central processing unit (CPU) and 0.3 gigabyte of computer memory in a traditional acoustic simulation of an underwater deposit. The image on the right used 240 CPUs running for up to 18 hours and 85 gigabytes of computer memory to generate a full-physics elastic simulation. The image contains a great deal more detail, including seismic S (shear) waves that travel in the earth.

atmosphere. However, some of the injected carbon dioxide would eventually leak into the atmosphere, where it would contribute to climate change.

To evaluate this approach to mitigating global warming, the Department of Energy formed a Center for Research on Ocean Carbon Sequestration, located at Lawrence Livermore and Lawrence Berkeley national laboratories. For one of the center's studies, Ken Caldeira, codirector of the center, along with colleagues Philip Duffy of Atmospheric Sciences and Mike Wickett of the Center for Applied Scientific Computing, used TC2K to evaluate the effectiveness of ocean carbon sequestration over a period of 100 years. "We want to know how much of the injected carbon dioxide would leak out of the ocean and at what rate," says Duffy. Their simulations showed that leakage into the atmosphere is much less when carbon dioxide injection is done at greater depths. The simulations also showed that maximum leakage may occur far from the injection site in areas where there is vigorous overturning of the ocean, such as in the North Atlantic. The simulations, the highest resolution of their kind ever done, used up to a hundred TC2K nodes and required a total of about 10,000 CPU hours.

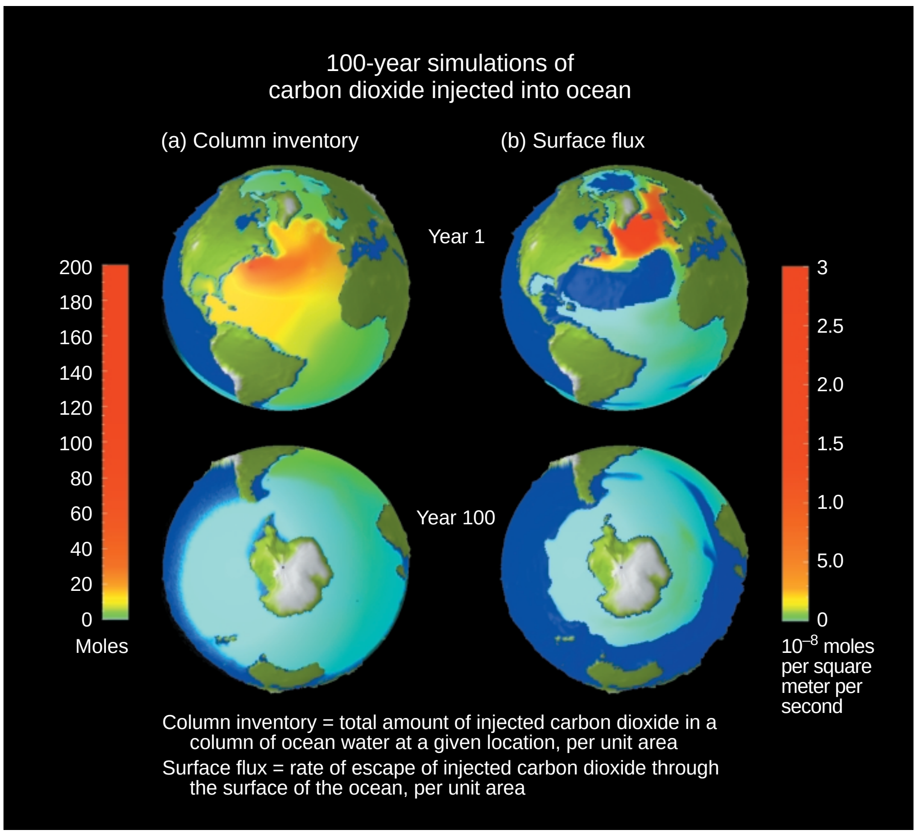

Multiprogrammatic and Institutional Computing supercomputers are making it possible to evaluate the potential effectiveness of injecting carbon dioxide into the ocean to mitigate global warming. These 100-year simulations depict what happens to carbon dioxide injected into the ocean off New York City at a depth of 710 meters. The column inventory, shown in the two spheres at the left, depicts the carbon dioxide traveling a great distance in 100 years. The surface flux, at right, shows that maximum escape of the injected carbon dioxide can occur far from the injection site in areas where there is vigorous overturning of the ocean, such as the North Atlantic.

Duffy and other atmospheric researchers, together with collaborators from Livermore's Center for Applied Scientific Computing, have used TC2K and other DOE supercomputers to perform the highest-resolution global climate simulations ever done. Duffy notes that as the simulation resolution gets finer, topographic features like mountains and valleys are represented better; these features have a direct bearing on weather. The simulations show that in some regions, maximum warming will occur in high-elevation regions because of a snow-albedo feedback: warming causes reduced snow cover, which in turn amplifies the warming by reflecting less sunlight back into space.

The simulations were performed in part on TC2K, in part at the National Energy Research Supercomputing Center at Lawrence Berkeley, and in part on ASCI Frost. "TC2K has really allowed us to push the limits of model resolution," says Duffy. "We're doing things that no other researchers can duplicate."

Unprecedented Ozone Studies

Atmospheric scientists in the Atmospheric Chemistry Group have used a hundred processors for a total of about 400 hours in their studies of the atmosphere. They have the only atmospheric chemistry model, IMPACT, that is capable of simulating the chemical reactions occurring in both the troposphere (the 10 kilometers of the atmosphere closest to Earth) and stratosphere. Past studies modeled the troposphere and stratosphere independently because of computational limitations. TC2K has enough computational power to permit coupling the troposphere and stratosphere in an atmospheric chemistry simulation capable of accurately predicting ozone concentrations. The results showed that studying interactions between these two regions of the atmosphere is important for understanding global and regional ozone distributions.

Understanding the ozone distribution throughout the atmosphere is crucial to the ability to predict not only the possibility of future stratospheric ozone depletion but also Earth's radiation balance and the magnitude of global warming. The scientists' model included resolutions of 2 degrees latitude by 2.5 degrees longitude, 46 levels of altitude from Earth's surface to 60 kilometers, tropospheric and stratospheric chemistry and physics involving 100 chemical species and 300 chemical reactions, and weather dynamics. "Because our codes are CPU-intensive, they do well on TC2K," says atmospheric scientist Doug Rotman. "The machine excels at big problems involving a lot of parameters." He says that while one goal is to keep increasing the resolution of the simulations, another goal is to include additional physics to make simulations more realistic.

Engineers David Clague, Elizabeth Wheeler, and Todd Weisgraber and University of California at Berkeley student Gary Hon have been using TC2K to perform three-dimensional simulations of both the stationary and mobile particles in portable microfluidic devices. These devices are being designed by Livermore researchers to automatically detect and identify viruses, bacteria, and toxic

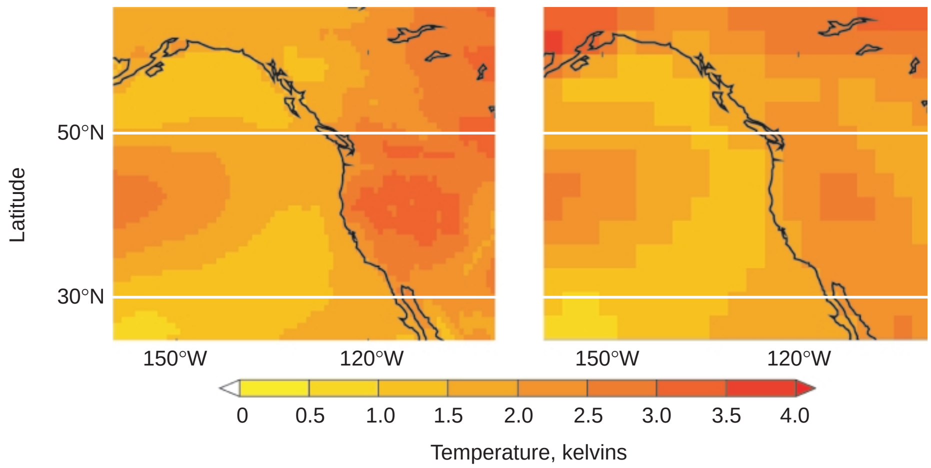

Simulations depict changes of surface air temperature between 2000 and 2100 in the western U.S. The simulation on the left, produced on Livermore's ASCI Frost and TC2K supercomputers, has a 75-kilometer resolution. The simulation on the right, produced on the National Energy Research Supercomputing Center at Lawrence Berkeley National Laboratory, has a 300-kilometer resolution. Its finer resolution gives a more detailed prediction that should be more accurate.

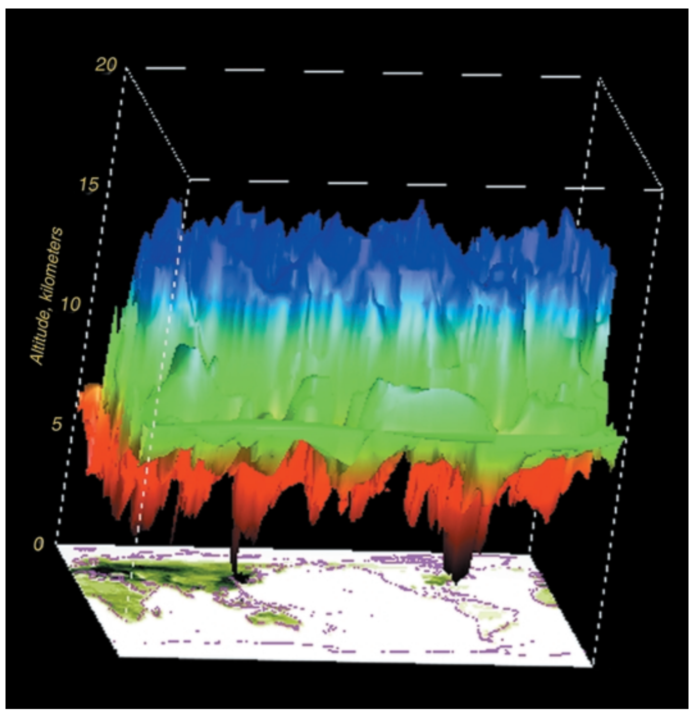

The TC2K supercomputer permits simulations that for the first time couple the troposphere and stratosphere in modeling the distribution of ozone. Modeling results show that in various locations, ozone in the lower stratosphere, which is typically in 100-parts-per-billion concentrations, is transported to near-surface altitudes.

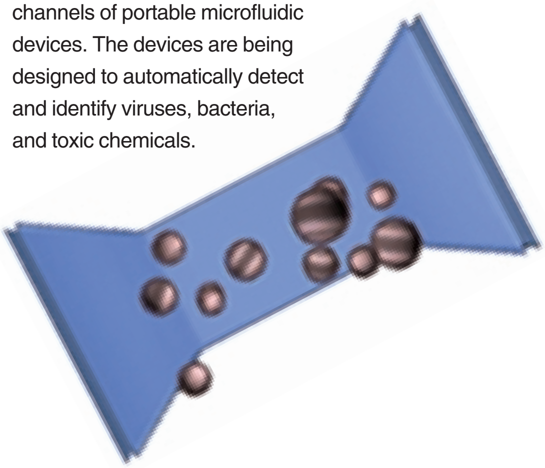

This simulation using TC2K shows the influence of electric fields on molecules traveling in narrow channels of portable microfluidic devices. The devices are being designed to automatically detect and identify viruses, bacteria, and toxic chemicals.

chemicals. (See *S&TR*, November 1999, pp. 10–16.) The devices have channels from 20 to 200 micrometers deep and up to a millimeter wide through which fluids travel. Because the channels are so small, intermolecular forces, which are typically masked in laboratory-scale instruments, affect the behavior of particles. The simulations show how beads and macromolecules are affected by each other's electric fields as they travel through a channel.

Initiative a Big Success

McCoy is pleased that M&IC computing resources have been so well received. One sign of the initiative's success has been the growing competition for the finite resources and the occasional wait of several days to begin big simulation runs.

"The desktop diaspora is over," says McCoy, and the result is unprecedented simulations and outstanding science. "We have achieved a balance in understanding what can be best done on the desktop and what can be best done in the experimental computational facility. We have developed close partnerships with our science teams, and we are already planning the next steps."

McCoy adds it is the "momentum based on continual change that keeps me and the Scientific Computing and Communications Department engaged and interested." That interest, he says, is based in part on the newest generation of simulations. They promise significant discoveries in science as Livermore researchers continue to elevate simulation to a level equal to that of theory and experiment.

—Arnie Heller

Key Words: Accelerated Strategic Computing Initiative (ASCI), ASCI Blue Pacific, ASCI Frost, ASCI White, atmospheric sciences, carbon sequestration, Compass Cluster, elastic simulations, E3D, Institutional Computing Executive Group, JEEP, microfluidic devices, Multiprogrammatic and Institutional Computing Initiative, parallel computing, Scientific Computing and Communications Department, Sunbert, supercomputers, TeraCluster, TeraCluster2000 (TC2K), Terascale Simulation Facility.

For further information contact Mike McCoy (925) 422-4021 (mccoy2@llnl.gov) or Greg Tomaschke (925) 423-0561 (tomaschke1@llnl.gov).

About the Scientist

MIKE McCOY, deputy associate director of Scientific Computing and Communications in Livermore's Computation Directorate, received his A.B. (1969) and Ph.D. (1975) in mathematics from the University of California at Berkeley. He joined the Laboratory in 1975 as a student employee. Upon completing his doctoral dissertation, he became a staff scientist in the National Energy Research Supercomputer Center, where he took responsibility for the development of algorithms for plasma codes. He went on to become group leader of the center's Massively Parallel Computing Group and then its deputy director. In the latter role, he directed the procurement of the 256-processor T3D computer, which was used for unclassified science. Now, as deputy associate director for Computation, he continues to support programs to advance network communications and security and to foster the development and integration of systems of powerful computers. He has worked with Livermore science teams to establish institutional computing to provide Laboratory scientists with access to powerful simulation environments. This sharing of institutional resources is part of his vision for enhancing the role of simulation in the scientific triad of theory, experiment, and simulation.